

ε -Nets for Halfspaces Revisited*

Sariel Har-Peled[†] Haim Kaplan[‡] Micha Sharir[§] Shakhar Smorodinsky[¶]

October 14, 2014

“It is a damn poor mind indeed which can’t think of at least two ways to spell any word.”

– Andrew Jackson

Abstract

Given a set P of n points in \mathbb{R}^3 , we show that, for any $\varepsilon > 0$, there exists an ε -net of P for halfspace ranges, of size $O(1/\varepsilon)$. We give five proofs of this result, which are arguably simpler than previous proofs [MSW90, CV07, PR08]. We also consider several related variants of this result, including the case of points and pseudo-disks in the plane.

1 Introduction

Since their introduction in 1987 by Haussler and Welzl [HW87] (see also Clarkson [Cla87] and Clarkson and Shor [CS89] for related concepts), ε -nets have become one of the central concepts in computational and combinatorial geometry, and have been used in a variety of applications, such as range searching, geometric partitions, and bounds on curve-point incidences, to name a few. We recall their definition: A *range space* (X, \mathcal{R}) is a pair consisting of an underlying universe X of objects, and a certain collection \mathcal{R} of subsets (*ranges*) of X . Of particular interest are range spaces of *finite VC-dimension*; skipping the exact definition, it suffices to require that, for any finite subset $P \subset X$, the number of distinct sets $r \cap P$, for $r \in \mathcal{R}$, is $O(|P|^d)$, for some constant d (which is upper bounded by the VC-dimension of (X, \mathcal{R})).

Given a range space (X, \mathcal{R}) , a finite subset $P \subset X$, and a parameter $0 < \varepsilon < 1$, an ε -*net* for P (and \mathcal{R}) is a subset $N \subseteq P$ with the property that any range $r \in \mathcal{R}$ with $|r \cap P| \geq \varepsilon|P|$ contains an element of N . In other words, N is a hitting set for all the “heavy” ranges.

*Work on this paper by Sariel Har-Peled was partially supported by an NSF CAREER award CCR-0132901. Work by Haim Kaplan was partially supported by Grant 975/06 from the Israel Science Fund. Work by Micha Sharir was partially supported by NSF Grant CCF-05-14079, by a grant from the U.S.-Israeli Binational Science Foundation, by grant 155/05 from the Israel Science Fund, Israeli Academy of Sciences, by a grant from the AFIRST joint French-Israeli program, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. Work by Shakhar Smorodinsky was partially supported by Grant 1136/12 from the Israel Science Fund.

[†]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu.

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; haimk@tau.ac.il

[§]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; michas@tau.ac.il

[¶]Department of Mathematics, Ben Gurion University of the Negev, Be’er Sheva 84105, Israel; shakhar@math.bgu.ac.il

The important result of Haussler and Welzl asserts that, for any (X, \mathcal{R}) , P , and ε as above, such that (X, \mathcal{R}) has finite VC-dimension d , there exists an ε -net N of size $O\left(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon}\right)$, and that a random sample of P of that size is an ε -net with constant probability. In particular, the size of N is independent of the size of P .

In geometric applications, this abstract framework is used as follows. The ground set X is typically a set of simple geometric objects (points, lines, hyperplanes), and the ranges in \mathcal{R} are defined as a Boolean combination of intersection with (or, for point objects, containment in) simply-shaped regions (halfspaces, balls, simplices, etc.), formally required to be regions of *constant descriptive complexity*, meaning that they are semi-algebraic sets defined in terms of a constant number of polynomial equations and inequalities of constant maximum degree. It is known that in such cases the resulting range space (X, \mathcal{R}) has finite VC-dimension (see, e.g., [Mat02]).

One of the major questions in the theory of ε -nets, posed since their introduction more than 25 years ago, is whether the factor $\log \frac{1}{\varepsilon}$ in the upper bound on their size is really necessary, especially in “normal” (geometric) situations. To be precise, it has been known, pretty early in the game, that in the general abstract context the answer is “yes”. This was shown by Komlós, Pach and Woeginger [KPW92] back in 1992, using a randomized construction on abstract hypergraphs (see also [PA95]).

Concerning simple geometric range spaces, Alon [Alo12] was the first to obtain a “non-linear” lower bound (i.e., more than C/ε for any positive constant C) on the size of ε -nets for the range space of points and triangles in the plane. More recently, Pach and Tardos [PT13] showed a lower bound of $\Omega(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$ for points and axis-parallel rectangles in the plane, and, more significantly, a lower bound $\Omega(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ for the dual range space (namely, a range space where X is a collection of axis-parallel rectangles, and each range is the subset of the rectangles that contains some given point). A simple reduction then leads to a lower bound of $\Omega(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$ on the size of ε -nets for points and halfspaces in \mathbb{R}^4 . In other words, ε -nets of size smaller than $\Theta(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$, even in simple geometric contexts, seem to be a relatively rare phenomenon.

Nevertheless, many improved upper bounds on the size of ε -nets in a variety of geometric contexts have been obtained [AES10, CV07, KK11, PR08]. Some of these bounds are linear (in $1/\varepsilon$), while others are in between linear and the general bound $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$. The first linear bound on the size of ε -nets has been obtained by Matoušek, Seidel and Welzl [MSW90], 25 years ago, for the special cases of points and halfspaces in two and three dimensions, and for some other related special cases. The proofs in [MSW90] are somewhat involved, and appear to have some technical difficulties, which are corrected in a revised version. Matoušek has given an alternative construction for halfspaces (in two and three dimensions), with the same bounds, using his shallow-cutting lemma [Mat92b]. Additional progress was made more recently. Clarkson and Varadarajan [CV07], essentially adapting Matoušek’s technique to their more general setting, have come up with a technique for constructing small-size ε -nets in geometric dual range spaces, where, as above, the objects of X are simply shaped regions, and each range is the subset of regions that are stabbed by some point. If the combinatorial complexity of the union of any finite number r of the regions is small, specifically $o(r \log r)$, then the corresponding range space admits ε -nets of size $o(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon})$. Pyrga and Ray [PR08] have proposed a general abstract scheme for constructing small-size ε -nets in hypergraphs (that is, range spaces) which satisfy certain properties, and have applied it to the special cases of halfspaces in two and three dimensions, and to a few other related instances. Later, Aronov, Ezra and Sharir [AES10] have shown the existence of ε -nets of size $O(\frac{1}{\varepsilon} \log \log \frac{1}{\varepsilon})$ for the range space of points and axis-parallel rectangles in the plane (a tight bound as subsequently shown by Pach and Tardos), as well as improved bounds for the size of ε -nets for several other range spaces. See also King and Kirkpatrick [KK11] for another recent improved bound for ε -nets for range spaces related to art-gallery visibility.

Our results. In this note we re-examine the techniques in [CV07, Mat92b, PR08], extract from them the main ingredients, and dress them up in different and arguably simpler proofs,^① which exploit the geometry of arrangements of planes in 3-space and of several other geometric structures. One of our (more complicated) proofs is essentially a restatement of Matoušek’s proof in [Mat92b], and of the similar (more general) proof in Clarkson and Varadarajan [CV07]; we give it here for the sake of completeness. The other four proofs are (in our opinion) considerably simpler. They use as a key ingredient a construction of Pyrga and Ray [PR08], but the analyses of the size of the resulting structure are different.

It is our hope that the multitude of proofs will shed more light on the problem, and that some of them might eventually be extended to other geometric range spaces.

Finally, we note that, for the case of halfspaces in three (and two) dimensions, one can construct linear-size ε -nets in an efficient deterministic manner, by a careful implementation of the construction given in the proofs, using standard techniques due to Matoušek [Mat96].

2 Small-size ε -nets for halfspaces in \mathbb{R}^3

Let P be a set of n points in \mathbb{R}^3 in general position, and let \mathcal{H} be the family of all (closed) halfspaces (bounded by planes). The main result of this section is:

Theorem 2.1 *Given a set P of n points in \mathbb{R}^3 in general position, and a parameter $0 < \varepsilon \leq 1$, there exists an ε -net for (P, \mathcal{H}) of size $O(1/\varepsilon)$.*

Remark. Clearly, the theorem implies that a similar result holds for the range space of points and halfplanes in the plane. This, however, can be established using considerably simpler and shorter proofs, see, e.g., [SY12, Woe89]. The general position assumption in the theorem is for the sake of simplicity of exposition, and was used also in the previous work [MSW90].

We give five different proofs of the theorem. The first four use the same construction, inspired by the approach of Pyrga and Ray [PR08], but differ in the way they show that the resulting ε -net has small size. The fifth proof is essentially a restatement of the proof of Matoušek [Mat92b], and of the similar more general proof in [CV07], and is included here for the sake of completeness; it is (in our opinion) somewhat more complex than the first four proofs.

The construction. Without loss of generality, it suffices to construct an ε -net for *lower* halfspaces. A symmetric construction will yield an ε -net for *upper* halfspaces, and the union of the two nets will be an ε -net for all halfspaces. Let \mathcal{H}^- denote the set of all lower halfspaces.

We fix some constant fraction $0 < \beta < 1$ (different choices of β will be made in the different proofs), and construct a maximal collection \mathcal{F} of lower halfspaces with the following properties (where we assume that $\varepsilon \leq 1/2$, for otherwise there exists a constant-size ε -net, by the general theory in [HW87]):

- (a) Each halfspace in \mathcal{F} contains between εn and $2\varepsilon n$ points of P .
- (b) For any pair of distinct halfspaces $h, g \in \mathcal{F}$, we have $|h \cap g \cap P| \leq \beta \varepsilon n$.

^①Obviously, simplicity is in the eye of the beholder, and we can only hope that the reader share our feeling that the proofs are indeed simpler.

It is easily seen that \mathcal{F} is finite.

For each halfspace $h \in \mathcal{F}$, we construct a $(\beta/2)$ -net N_h for the set system $(h \cap P, \mathcal{H}^-)$, of size $O\left(\frac{1}{\beta} \log \frac{1}{\beta}\right) = O(1)$, using the standard bounds on the size of ε -nets [HW87], and form the union $N^{(1)} = \bigcup_{h \in \mathcal{F}} N_h$.

We next repeat the same construction for each value $\varepsilon_j = 2^{j-1}\varepsilon$, for $j = 1, 2, \dots$, using the same parameter β for each ε_j . We obtain a sequence of subsets $N^{(j)} \subseteq P$, and we set N to be their union. Let us also denote by $\mathcal{F}^{(j)}$ the maximal set of halfspaces constructed at the j th step.

N is an ε -net. It is easy to see that N is an ε -net for (P, \mathcal{H}^-) . Indeed, let h be a lower halfspace which contains at least εn points of P . There exists $j \geq 1$ such that $2^{j-1}\varepsilon n \leq |h \cap P| < 2^j \varepsilon n$. If $h \in \mathcal{F}^{(j)}$ then it certainly contains a point of N (it contains the nonempty subset $N_h \subseteq N$). Otherwise, $\mathcal{F}^{(j)} \cup \{h\}$ must violate property (b), so $\mathcal{F}^{(j)}$ contains a lower halfspace g such that

$$|h \cap g \cap P| > \beta \varepsilon_j n = \beta 2^{j-1} \varepsilon n \geq \frac{\beta}{2} |g \cap P|.$$

Hence, by construction, h must contain a point of N_g , and thus of N .

The main challenge is to argue that $|\mathcal{F}| = |\mathcal{F}^{(1)}| = O(1/\varepsilon)$. Since β is a constant, this would imply that $|N^{(1)}| = O(1/\varepsilon)$ too, and, more generally, that $|N^{(j)}| = O(1/(2^{j-1}\varepsilon))$. Summing these bounds, we would then obtain $|N| = O(1/\varepsilon)$.

Remark. The approach just presented, which uses a geometric progression of values of ε and a corresponding sequence of constructions, appears to be unnecessarily over-complicated. Specifically, for halfspaces, a single step suffices: We only construct $N^{(1)}$, and claim that it is an ε -net. Indeed, if h is any lower halfspace containing at least εn points of P then, by the general position assumption, we can shrink it to another halfspace $h' \subseteq h$ which contains exactly εn points of P . The preceding argument implies that h' contains a point of $N^{(1)}$, and therefore so does h . The reason for complicating the construction is that it can also handle ranges which do not have this “shrinking property”, namely the property that any range can be shrunk to a smaller range which contains any prescribed number of points of P . The same trick of repeatedly doubling ε has also been used by Pyrga and Ray [PR08].

We give four proofs of the claim that $|\mathcal{F}| = O(1/\varepsilon)$. Each proof is based on a specific (and different) choice of β , which is spelled out during the analysis.

First proof. For each $h \in \mathcal{F}$ let π_h denote its bounding plane. By slightly perturbing these planes, without changing any of the subsets $h \cap P$, for $h \in \mathcal{F}$, we may assume that the planes π_h are in general position. We claim that all the planes π_h appear on their upper envelope E . Indeed, suppose to the contrary that there exists $h \in \mathcal{F}$ such that π_h lies fully below the envelope. Let v be the vertex of the envelope closest to h . Clearly, the union of the three halfspaces $h_1, h_2, h_3 \in \mathcal{F}$ defining v cover h ; that is, $h \subseteq h_1 \cup h_2 \cup h_3$. Hence, for at least one index $i \in \{1, 2, 3\}$, we have $|h \cap h_i \cap P| \geq \frac{1}{3} |h \cap P| \geq \frac{1}{3} \varepsilon n$, which contradicts property (b) if we choose $\beta < 1/3$.

Put $t = |\mathcal{F}|$, and consider E as a planar map, which has t faces. Define the *degree* $\deg(f)$ of a face f of E , lying on some plane π_h , to be the number of planes π_g which appear on the 1-level of $\mathcal{A}(\mathcal{F})$ directly below f (see Figure 1 for an illustration). In general, each such plane π_g either meets ∂f or contributes a face to the 1-level which lies fully below f ; the second case is impossible, though, for then π_g would not appear on the upper envelope. Hence, assuming general position, $\deg(f)$ is equal to the number of edges of f .




Figure 1: A pocket.

By Euler's formula, the number of edges of the upper envelope of t planes in \mathbb{R}^3 is at most $3t - 6$. Since each such edge participates in two pockets, we get $\sum_f \deg(f) < 6t$, where the sum extends over all faces f of E . Hence, at least half of the t faces of E have degree at most 11; refer to these faces as *light*.

Let f be one of these faces, and let h be the corresponding halfspace. Then we have

$$h \setminus (\cup \mathcal{F} \setminus \{h\}) = h \setminus \cup \mathcal{F}_h,$$

where \mathcal{F}_h is the set of all halfspaces which contribute to the degree of f ; refer to this expression as the *pocket* of f (or of h). By the choice of f , we have $|\mathcal{F}_h| \leq 11$. Hence, the number of points of P in the pocket is at least

$$|h \cap P| - \sum_{g \in \mathcal{F}_h} |g \cap h \cap P| \geq \varepsilon n - 11\beta\varepsilon n \geq \frac{1}{2}\varepsilon n,$$

if we choose $\beta \leq 1/22$.

Since the pockets are clearly pairwise disjoint, the overall number of points of P in the pockets of the at least $t/2$ light faces is at least $\frac{1}{4}t\varepsilon n$. Hence we have $\frac{1}{4}t\varepsilon n \leq n$, implying that $t \leq 4/\varepsilon$, as claimed. \square

Second proof. Here, in an attempt to generalize the analysis to other kinds of range spaces, we do not assume that all the bounding planes appear on the upper envelope. Let \mathcal{F}_0 denote the subset of those halfspaces in \mathcal{F} whose bounding planes do appear on the envelope of \mathcal{F} , and put $t_0 = |\mathcal{F}_0|$. For $i \geq 1$, let \mathcal{F}_i be the subset of halfspaces in $\mathcal{G}_i = \mathcal{F} \setminus (\mathcal{F}_0 \cup \dots \cup \mathcal{F}_{i-1})$ whose bounding planes appear on the upper envelope of \mathcal{G}_i . We repeat this peeling process, so at the last stage k , all of \mathcal{F} is exhausted, and $\mathcal{F}_{k+1} = \emptyset$.

The overall complexity of the 0- and 1-levels of $\mathcal{A}(\mathcal{G}_i)$ is at most $c(t_i + t_{i+1})$, where $t_i = |\mathcal{F}_i|$, and where c is some absolute constant. This follows from Euler's formula for the 0-level (as in the first proof), combined with the random sampling technique of Clarkson and Shor [CS89] (or, alternatively, the simpler scheme of Tagansky [Tag96]). For this, note that any plane that shows up on the 0- or 1-level in $\mathcal{A}(\mathcal{G}_i)$ must belong to $\mathcal{F}_i \cup \mathcal{F}_{i+1}$. That is, every time we peel away the halfspaces touching the envelope, we also remove the pockets touching the envelope, "exposing" their lower boundaries to the upper unbounded cell, and making these boundaries participate in the next envelope. Recall that here we no longer assume that all the planes bounding a pocket reach the boundary of its upper facet, as in the preceding proof. After the removal of \mathcal{F}_i , the new pockets of the bounding planes of halfspaces in \mathcal{F}_{i+1} are (pairwise disjoint and) disjoint from the pockets of the preceding stages.

Summing over all stages, we obtain a total of $t = t_0 + t_1 + \dots + t_k$ pockets (each plane contributes exactly one pocket), and the sum of the degrees of the corresponding halfspaces (which we bound by the

complexity of their pockets) is at most $c((t_0 + t_1) + (t_1 + t_2) + \dots + (t_{k-1} + t_k) + (t_k)) \leq 2ct$. Hence, the average degree of a halfspace is at most $2c$, and therefore at least half of the pockets are of halfspaces with degree at most $4c$. Hence, choosing $\beta = 1/(8c)$, and arguing as above, the bound $t = O(1/\varepsilon)$ follows. \square

Third proof. The preceding proof peels off the halfspaces of \mathcal{F} in batches, where in each step we remove many halfspaces. An alternative, and perhaps somewhat more natural approach is to peel them off one by one. In this proof we use a random peeling order, and exploit known properties of randomized incremental constructions of upper envelopes to establish the desired bound on the size of \mathcal{F} .

Specifically, draw a random permutation of \mathcal{F} , which we write as (h_1, h_2, \dots, h_t) , and insert the halfspaces one by one in this order. When h_j is inserted, it adds to the union a new (possibly empty) pocket, which, as above, is $h_i \setminus \bigcup_{j < i} h_j$; the pocket is nonempty if and only if π_{h_j} appears on the current upper envelope E_j . An obvious but crucial property is that all these pockets are pairwise openly disjoint. To facilitate the analysis, we maintain a triangulation of E_j , using, e.g., bottom-vertex triangulation of each facet, and maintain for each triangle τ a *conflict list* of all the halfspaces h not yet inserted, which meet τ . If τ lies fully in the interior of such a halfspace h , then adding h removes τ completely from the envelope; otherwise, adding h splits τ into a portion which is hidden from the envelope and a portion which remains on the envelope. The remaining portions of the envelope are re-triangulated, the conflict lists of the new triangles are computed from the lists of the destroyed triangles, and the process continues. See, e.g., Seidel [Sei93] for a review of randomized incremental constructions of this kind.

The standard analysis of randomized incremental constructions (see [Sei93]) implies that the expected overall number of triangles that are ever created by the algorithm is at most ct , for some absolute constant $c > 0$.

Note that the complexity of the pocket of h_j , and thus the degree of h_j , *at the time of its insertion*, is proportional to the number of triangles of E_{j-1} that are killed by h_j , either by being fully eliminated, or by being split and replaced by new triangles; refer to these degrees as the *degrees at birth*. It follows that the expected sum of the degrees at birth of the halfspaces is at most $c't$, for another absolute constant $c' > 0$. Hence the average degree at birth is at most c' , so at least half of the halfspaces have degree at birth at most $2c'$.

Choose $\beta = 1/(4c')$. Arguing as above, it is easy to see that the number of points of P in the pocket at birth of such a ‘light’ halfspace is at least $\frac{1}{2}\varepsilon n$. Hence the total number of points of P in these pockets is at least $\frac{1}{4}t\varepsilon n$. Since the pockets at birth are pairwise disjoint, the bound on t follows. \square

Fourth proof. This time we pass to the dual space, where each point $p \in P$ is mapped to a dual plane p^* , and each range $h \in \mathcal{F}$ is mapped to a dual point h^* (actually, dual to the plane bounding h), so that point p lies in halfspace h if and only if the dual point h^* lies above the dual plane p^* . Let P^* (resp., \mathcal{F}^*) denote the resulting set of n dual planes (resp., of t dual points). Note that the level of each point $h^* \in \mathcal{F}^*$ in the arrangement $\mathcal{A}(P^*)$ is between $k = \varepsilon n$ and $2k$, and that, for any pair of distinct points $h^*, g^* \in \mathcal{F}^*$, the number of planes that separate them is at least $2(1 - \beta)k$. Indeed, each of these planes passes below exactly one of h^*, g^* . Since at least k planes pass below h^* , at least k planes pass below g^* , and at most βk planes pass below both of them, the claim follows. Choose $r = (1 - \beta)k$, and define, for each dual point $h^* \in \mathcal{F}^*$, the “ball” B_h of all the vertices of $\mathcal{A}(P^*)$ which are separated from h^* by fewer than r planes. As argued by Welzl [Wel92], the number of vertices in B_h is $\Omega(r^3) = \Omega(k^3)$. By construction, no vertex of $\mathcal{A}(P^*)$ can appear in more than one of these balls. Indeed, let $d_{P^*}(u, v)$ denote the number of planes of P^* which separate u and v ; then d_{P^*} satisfies the triangle inequality, from which the claim follows readily. Moreover, since h^* lies at level between k and $2k$, the level of any

of these vertices is between $k - r$ and $2k + r$, so they all lie at level at most $3k$. As shown by Clarkson and Shor [CS89], the overall number of vertices of $\mathcal{A}(P^*)$ at level at most $3k$ is $O(nk^2)$. Hence the number t of balls satisfies $tk^3 = O(nk^2)$, and thus $t = O(n/k) = O(1/\varepsilon)$. \square

Fifth proof. This proof is an adaptation of Matoušek’s proof on shallow cutting [Mat92b, Section 5], and also of the proof in [CV07], which itself is an extension of Matoušek’s proof; we include it here for the sake of completeness. In this proof we abandon the set \mathcal{F} . As in the preceding proof, we dualize the points of P to planes, and denote the resulting set of planes by P^* .

Choose a random sample N_0 of $r = \lceil 2/\varepsilon \rceil$ points of P , pass to the subset $N_0^* \subseteq P^*$ of dual planes, construct their lower envelope E_0 , and triangulate each of its faces, into a total of $O(r)$ triangles. We extend each of these triangles τ_0 downwards to the semi-unbounded vertical prism τ bounded from above by τ_0 . We thus obtain a decomposition of the region E_0^- below E_0 into $O(r)$ such prisms.

We take N_0 to be part of the output ε -net. By construction, any lower halfspace h whose bounding plane is dualized into a point h^* which lies above E_0 has a plane of N_0^* passing below it, so, in primal space, h contains the point of N_0 dual to that plane. Thus it remains to consider halfspaces whose dual points lie below E_0 .

By the ε -net theory, our sample N_0 satisfies, with high probability, the property that each of the prisms τ in the decomposition of E_0^- is crossed by at most $\frac{cn}{r} \log r$ planes of P^* , for some absolute constant c . We may assume that our sample does indeed satisfy this property.

Let τ be one of the prisms, and suppose that τ is crossed by tn/r planes of P^* . If $t < 1$, we leave τ intact. Otherwise, we take a random sample N_τ^* of $c't \log t$ planes from those crossing τ , for some sufficiently large but absolute constant c' . We may assume that each of these samples is a $(1/t)$ -net, for the corresponding value of t , for the range space of the planes of P^* crossing τ , where each range is the subset of planes stabbed by some downward-directed vertical ray (that is, lying below some point).

We take N to be the union of N_0 and of all the primal sets N_τ (dual to the corresponding N_τ^*), and claim that N is an ε -net. That is, if h is a lower halfspace that contains at least εn points of P then h contains a point of N . Indeed, as already noted, it suffices to consider the case where the point h^* dual to the plane bounding h lies in E_0^- . By the same reasoning, if h^* lies in some initial prism τ and above the upper envelope E_τ of N_τ^* , we are also done. The remaining case is impossible, because then the downward-directed vertical ray emanating from h^* does not meet any plane of N_τ^* , and therefore meets fewer than εn planes of P^* , contradicting the assumption that h is “heavy”.

It remains to show that $|N| = O(r)$. This follows from the exponential decay lemma, initially established by Chazelle and Friedman [CF90], and later extended by Agarwal *et al.* [AMS98]. (This lemma holds with high probability for the initial sample N_0 , and, as above, we may assume that the lemma does indeed hold for N_0 .)

This completes the fifth proof of the theorem. \square

Remarks. (1) It is interesting to compare the construction used in the first four proofs to the one used in the fifth. A common feature of both constructions is that their last step involves the construction of ε' -nets, for ε' a *constant*,^② for many subsets of P . However, these subsets are obtained by what seems to be totally unrelated methods. It would be interesting to “see through” these constructions, and perhaps show that they are more related than what meets the eye.

(2) It is worth mentioning that the construction of Pyrga and Ray, on which our proofs are based, might be very inefficient in some simple geometric settings (other than the one considered here). For example, consider the range space defined by points and lines in the plane. We use the construction of Elekes

^②Well, in the fifth proof, constant on average.

[Ele02] of n points and n lines in the plane with $\Omega(n^{4/3})$ incidences. Let k be an integer and put $n = 2k^3$. Consider the set P consisting of the points in the integer grid $[1 : k] \times [1 : 2k^2]$. Put $\varepsilon = \frac{1}{2k^2}$. Notice that each of the lines of the form $y = ax + b$, for $a \in [1 : k]$ and $b \in [1 : k^2]$ contains exactly $k = \varepsilon n$ points of P . Moreover, every pair of such lines intersect in at most one point of P . Hence, the collection \mathcal{F} of these lines satisfies conditions (a) and (b) of the construction, but $|\mathcal{F}|$ is way too large, namely $|\mathcal{F}| = k^3 = \Omega(\frac{1}{\varepsilon^{3/2}})$.

2.1 Deterministic construction

In this final subsection, we note that, for the case of halfspaces in three (and two) dimensions, one can construct linear-size ε -nets in an efficient deterministic manner, by a careful implementation of the construction given in the first four proofs of Theorem 2.1, using standard techniques due to Matoušek and others. Specifically, we will only consider the problem in three dimensions, since the two-dimensional case can be handled as a special case. Moreover, a different deterministic construction in the plane follows from the technique of Matoušek in [Mat90].

We proceed as follows. Consider the dual version, in which we have a collection H of n planes in \mathbb{R}^3 , and the ranges correspond to points, so that the range of a point q is the set of planes passing below q . Consider a modified range space, in which the ranges are defined by segments, so that the range of segment e is the set of planes of H which cross e .

Set $\varepsilon' = \varepsilon/4$. We first construct an ε' -approximation X , of size $|X| = O(1/\varepsilon^2 \log(1/\varepsilon))$, for the modified range space. That is, for each segment e , if we set H_e (resp., X_e) to be the set of those planes in H (resp., X) that cross e , then we have

$$\left| \frac{|X_e|}{|X|} - \frac{|H_e|}{|H|} \right| < \frac{\varepsilon}{4}.$$

One can construct X deterministically in time $O(n/\varepsilon^c)$, where c is some (small) constant, using the general technique of Matoušek [Mat96].

Next, we construct the arrangement $\mathcal{A}(X)$, and extract from it the set S of all the vertices at level at most $\frac{3}{2}\varepsilon|X|$. Our goal is to compute a maximal subset $\mathcal{F} \subseteq S$, such that the crossing distance between any pair of vertices in \mathcal{F} is at least $r := \frac{1}{4}\varepsilon|X|$. We do this by explicitly computing the crossing distance (in $\mathcal{A}(S)$) between each pair of points of S , and then by augmenting \mathcal{F} incrementally. We first set \mathcal{F} to be the empty set. At each step we take the next point z in S , and check whether it lies at crossing distance larger than r from all the points currently in \mathcal{F} . If so, we add it to \mathcal{F} . It is easy to see (using the fourth proof of Theorem 2.1) that $|\mathcal{F}| = O(1/\varepsilon)$.

Next, for each point $p \in \mathcal{F}$, we compute a β -net N_p , for some $\beta < 1/8$, for the set H_p of all the planes of H lying below p . This takes $O(\varepsilon n/\beta^c)$ time per point, using the algorithm in [Mat96], for an overall time $O(n/\beta^c) = O(n)$. We take N to be the union of all the sets N_p , and claim that (the primal image of) N is indeed an ε -net for (P, \mathcal{H}^-) . Since $|N| = O(|\mathcal{F}|) = O(1/\varepsilon)$, we have indeed constructed an ε -net for (P, \mathcal{H}^-) of size $O(1/\varepsilon)$.

The argument that N is an ε -net proceeds as follows. Denote the level of a point u in $\mathcal{A}(H)$ by λ_u . Let q be a point at level $\lambda_q = \varepsilon n$ in $\mathcal{A}(H)$. Letting e denote the downward-directed ray from q , we have $|H_e| = \varepsilon n$, and therefore

$$\frac{3\varepsilon}{4} < \frac{|X_e|}{|X|} < \frac{5\varepsilon}{4}.$$

Thus q lies at level between $\frac{3}{4}\varepsilon|X|$ and $\frac{5}{4}\varepsilon|X|$ in $\mathcal{A}(X)$. Let q' be a vertex of the cell in $\mathcal{A}(X)$ containing q . Then q' lies at the same level (with respect to X), and in particular $q' \in S$. This implies, as above,

that $\lambda_{q'} \leq \frac{3}{2}\varepsilon n$. Since $|X_{qq'}| = 0$, we have $|H_{qq'}| < \frac{1}{4}\varepsilon n$. Hence at least $\frac{3}{4}\varepsilon n$ of the planes of H passing below q also pass below q' . In other words, at least half of the $\lambda_{q'}$ planes of H passing below q' also pass below q .

Now if $q' \in \mathcal{F}$, then the β -net $N_{q'}$ contains a plane passing below q (by our choice, $\beta < 1/2$), and we are done. Otherwise, there exists a point $q'' \in \mathcal{F}$ such that the crossing distance $|X_{q'q''}|$ satisfies $|X_{q'q''}| \leq \frac{1}{4}\varepsilon|X|$, so $|H_{q'q''}| \leq \frac{1}{2}\varepsilon n$, and therefore

$$|H_{qq''}| \leq |H_{qq'}| + |H_{q'q''}| \leq \frac{3}{4}\varepsilon n.$$

Hence, at least $\frac{1}{4}\varepsilon n$ of the planes of H below q also pass below q'' . Observe that the level of q'' in $\mathcal{A}(H)$ is at most $(\frac{3}{2} + \frac{1}{4})\varepsilon n < 2\varepsilon n$, so at least $1/8$ of the planes below q'' also pass below q , so one of the planes of $N_{q''}$ must pass below q .

The resulting algorithm runs in (deterministic) time $O(n/\varepsilon^c)$, where c is some small constant. In this sketchy solution we made no attempt to optimize the dependence of the running time on ε , which can probably be improved using efficient algorithms for constructing cuttings [Mat92a].

References

- [AES10] B. Aronov, E. Ezra, and M. Sharir. Small-size ε -nets for axis-parallel rectangles and boxes. *SIAM J. Comput.*, 39(7):3248–3282, 2010.
- [Alo12] N. Alon. A non-linear lower bound for planar epsilon-nets. *Discrete Comput. Geom.*, 47(2):235–244, 2012.
- [AMS98] P. K. Agarwal, J. Matoušek, and O. Schwarzkopf. Computing many faces in arrangements of lines and segments. *SIAM J. Comput.*, 27(2):491–505, 1998.
- [CF90] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.
- [Cla87] K. L. Clarkson. New applications of random sampling in computational geometry. *Discrete Comput. Geom.*, 2:195–222, 1987.
- [CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [CV07] K. L. Clarkson and K. R. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete Comput. Geom.*, 37(1):43–58, 2007.
- [Ele02] G. Elekes. Sums versus products in number theory, algebra and Erdős geometry — a survey. In *Paul Erdős and his Mathematics II*, pages 241–290. Bolyai Math. Soc., 2002.
- [HW87] D. Haussler and E. Welzl. ε -nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
- [KK11] J. King and D. Kirkpatrick. Improved approximation for guarding simple galleries from the perimeter. *Discrete Comput. Geom.*, 46(2):252–269, 2011.
- [KPW92] J. Komlós, J. Pach, and G. Woeginger. Almost tight bounds for ε -nets. *Discrete Comput. Geom.*, 7:163–173, 1992.

- [Mat90] J. Matoušek. Construction of ε -nets. *Discrete Comput. Geom.*, 5:427–448, 1990.
- [Mat92a] J. Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8:315–334, 1992.
- [Mat92b] J. Matoušek. Reporting points in halfspaces. *Comput. Geom. Theory Appl.*, 2(3):169–186, 1992.
- [Mat96] J. Matoušek. Derandomization in computational geometry. *J. Algorithms*, 20:545–580, 1996.
- [Mat02] J. Matoušek. *Lectures on Discrete Geometry*, volume 212 of *Grad. Text in Math.* Springer, 2002.
- [MSW90] J. Matoušek, R. Seidel, and E. Welzl. How to net a lot with little: Small ε -nets for disks and halfspaces. In *Proc. 6th Annu. Sympos. Comput. Geom. (SoCG)*, pages 16–22, 1990. A corrected version is available at <http://kam.mff.cuni.cz/~matousek/enets3.ps.gz>.
- [PA95] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. John Wiley & Sons, 1995.
- [PR08] E. Pyrga and S. Ray. New existence proofs for ε -nets. In *Proc. 24th Annu. Sympos. Comput. Geom. (SoCG)*, pages 199–207, 2008.
- [PT13] J. Pach and G. Tardos. Tight lower bounds for the size of epsilon-nets. *J. Amer. Math. Soc.*, 26:645–658, 2013.
- [Sei93] R. Seidel. Backwards analysis of randomized geometric algorithms. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms and Combinatorics*, pages 37–68. Springer-Verlag, 1993.
- [SY12] S. Smorodinsky and Y. Yuditsky. Polychromatic coloring for half-planes. *J. Comb. Theory, Ser. A*, 119(1):146–154, 2012.
- [Tag96] B. Tagansky. A new technique for analyzing substructures in arrangements of piecewise linear surfaces. *Discrete Comput. Geom.*, 16(4):455–479, 1996.
- [Wel92] E. Welzl. On spanning trees with low crossing numbers. In *Data Structures and Efficient Algorithms, Final Report on the DFG Special Joint Initiative*, volume 594 of *Lect. Notes in Comp. Sci.*, pages 233–249. Springer-Verlag, 1992.
- [Woe89] G. Woeginger. Epsilon-nets for halfplanes. In *Proc. 14th Internat. Workshop Graph-Theoret. Concepts Comput. Sci.*, volume 344 of *Lect. Notes in Comp. Sci.*, pages 243–252, 1989.